Robust Unsupervised StyleGAN Image Restoration

Yohan Poirier-Ginter,[1,2] Jean-François Lalonde[1]

1 Université Laval, 2 Inria

JUNE 18-22, 2023 CVPR VANCOUVER, CANADA

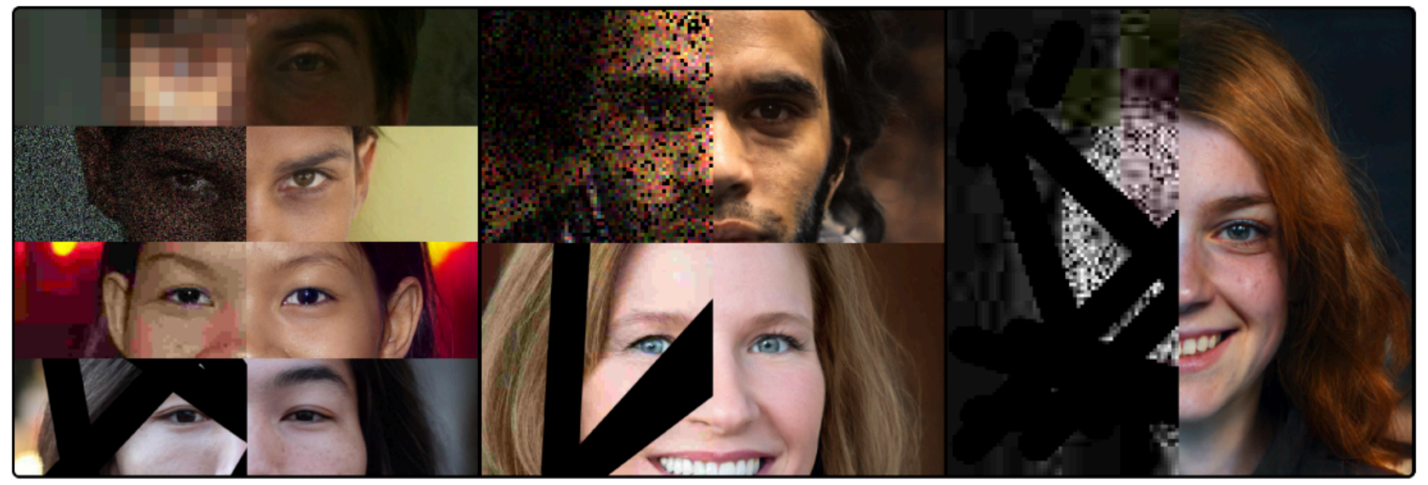

GOAL

- Our objective is to restore images corrupted by **known degradations,** without any training (fully **unsupervised**).
- Many existing approaches (PULSE, L-BRGM) exploit a pretrained StyleGAN generator:

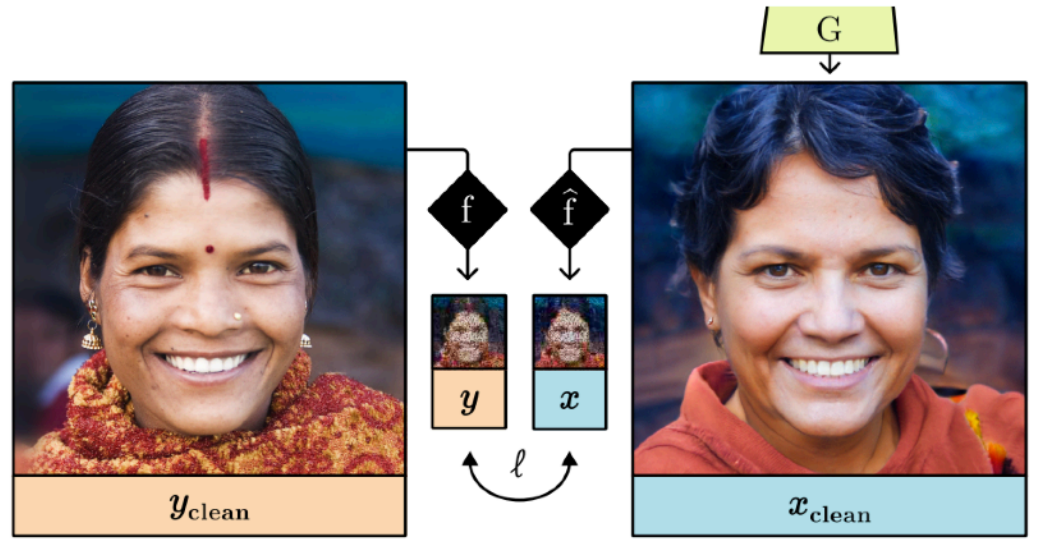

Using gradient descent, these methods **invert the generative process** to find an image which matches the target once degraded in the same way.

MOTIVATION

- They introduce regularization hyperparameters which tradeoff between quality of fit and realism.
- These hyperparameters are **not robust**, and must be adjusted for each task:

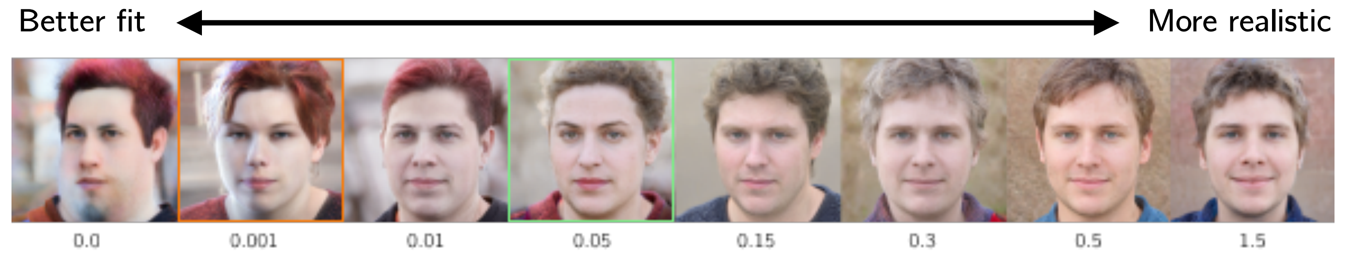

In this work, we propose a method which is **robust.**

METHOD

- In this work, we argue that such regularizers are <u>not actually required.</u>
- Instead, we propose combining:

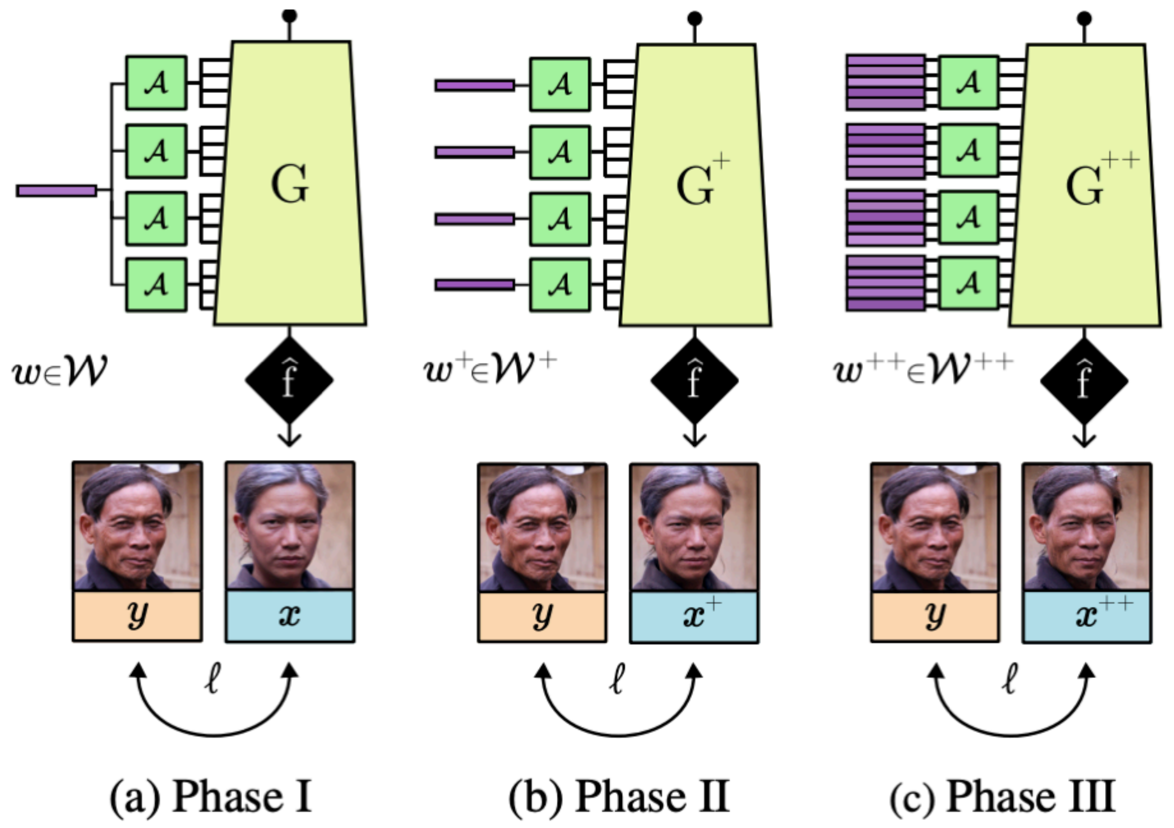

1) **A conservative optimizer** (normalized gradient decent)

Most artifacts are introduced by the Adam optimizer

2) **A progressive** (three-phase) **latent extension**

"Initialization is the best regularization"

Additionally, we use **3) a multiscale LPIPS** loss.

<u>Previous works use regularizers to remove artifacts introduced by the optimization process, while our careful optimization avoids introducing them in the first place</u>

SINGLE DEGRADATIONS

We compare our method to PULSE and L-BRGM on upsampling, inpainting, denoising, and deartifacting, at five different levels of degradation (extra-small "XS" to extra-large "XL")

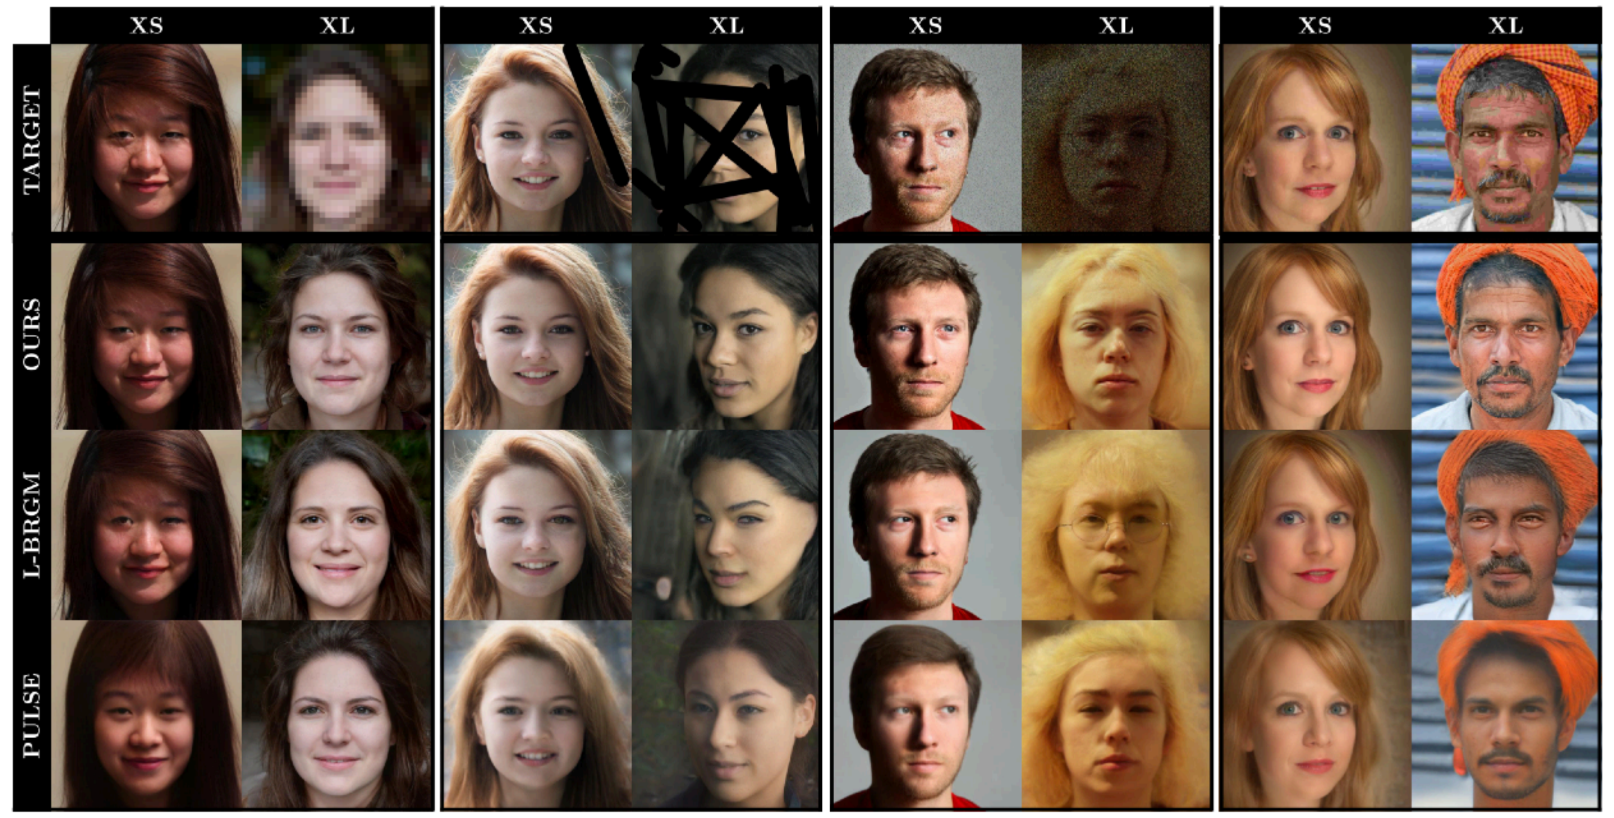

	Accur. (LPIPS) ↓			Fidelity (LPIPS) ↓			Realism (pFID) ↓		
	PULS	**L-BRG**	**OURS**	**PULS**	**L-BRG**	**OURS**	**PULS**	**L-BRG**	**OURS**
Upsampling (bilinear, bicubic or Lanczos)									
XS	.493	.407	.414	.432	.295	.313	44.5	23.6	17.0
S	.492	.412	.449	.353	.140	.239	34.3	25.5	22.0
M	.495	.458	.472	.261	.124	.172	29.3	35.4	22.3
L	.501	.487	.490	.185	.129	.127	21.9	26.0	20.9
XL	.512	.506	.514	.083	.095	.090	24.9	21.3	21.3
Denoising (clamped Poisson and Bernoulli mixture)									
XS	.501	.440	.425	.275	.152	.156	56.1	27.2	18.5
S	.499	.450	.434	.252	.138	.140	53.7	28.6	19.1
M	.500	.465	.446	.224	.155	.130	54.5	22.1	19.8
L	.501	.481	.457	.185	.138	.110	56.4	24.6	19.2
XL	.504	.511	.474	.134	.110	.084	49.4	25.1	17.9
Deartifacting (JPEG compression)									
XS	.498	.442	.432	.404	.341	.349	52.3	26.3	14.8
S	.497	.448	.437	.398	.352	.350	49.6	22.4	15.4
M	.498	.461	.445	.413	.357	.357	33.2	24.1	15.4
L	.500	.475	.460	.395	.367	.374	46.9	25.2	16.0
XL	.508	.503	.490	.427	.418	.412	30.8	22.1	18.7
Inpainting (random strokes)									
XS	.498	.409	.378	.464	.374	.348	46.9	24.4	12.9
S	.501	.425	.387	.356	.287	.264	42.3	27.2	14.2
M	.509	.438	.396	.283	.227	.206	38.5	30.1	14.5
L	.513	.452	.409	.231	.184	.163	32.6	33.1	15.3
XL	.524	.460	.422	.187	.157	.132	36.2	25.2	15.9

- For both baselines, we perform a hyper-parameter search for the best possible accuracies.
- For our method, we use the same hyperparameters found on a validation set.

COMPOSED DEGRADATIONS

Without hyperparameters to adjust, working solutions for different degradations can be combined without any change:

DISCUSSION

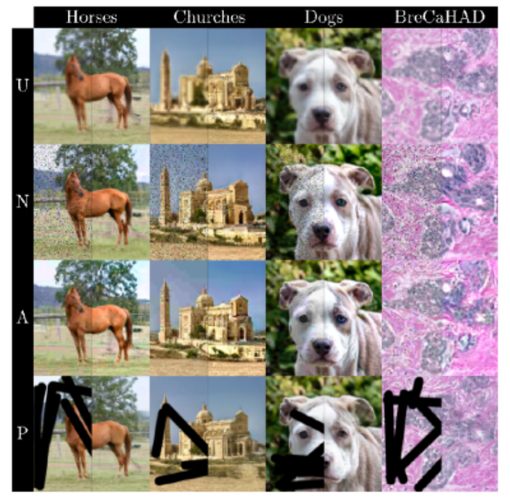

- We can easily use the same approach for different datasets.
- But, can StyleGAN be trained with a general-purpose purpose dataset? (See GigaGAN)
- Can the degradations models be inferred?

website: <u>https://lvsn.github.io/RobustUnsupervised/</u>

Acknowledgements: This research was supported by NSERC grant RGPIN-2020-04799, a MITACS Globalink internship to Y. Poirier-Ginter, and by the Digital Research, Alliance Canada.